

Crisis Action Planning and Replanning using SIPE-2

Jennifer D. Skidmore
 Rome Laboratory
 RL/C3CA
 Griffiss AFB, NY 13441-4505
 skiddles@aivax.rl.af.mil

Abstract

Rome Laboratory and DARPA are jointly sponsoring an initiative to develop the next generation of AI planning and scheduling technology focused on military operations planning, especially for crisis situations. SRI International has demonstrated their knowledge-based planning technology in this domain with a system called SOCAP, System for Operations Crisis Action Planning. SOCAP's underlying power comes from SIPE-2, a hierarchical, domain-independent, non-linear AI planner also developed at SRI. This paper discusses the features of SIPE-2 that made it an ideal choice for military operations planning, and which contributed greatly to SOCAP's success.

1 Introduction

The goal of the DARPA/Rome Lab Planning Initiative is to develop the next generation of artificial intelligence planning and scheduling tools focused on military operations planning, especially crisis action planning. Developing an appropriate course of action (COA) in response to a crisis situation encompasses the planning of both the *employment* of forces against the enemy and the *deployment* of forces and cargo to their required destinations in time to complete their mission.

This planning is done at various levels along the chain of command. Specifically, when a crisis occurs that requires military action, the Chairman of the Joint Chiefs of Staff passes down a mission statement with some planning guidance to the Commander-in-Chief (CINC) of the appropriate geographical area. The CINC and his staff are responsible for generating alternative COAs based on

available intelligence data about the enemy's posture, terrain information, approaches that have proven successful in the past, logistics capabilities, and many other factors. The CINC is mostly concerned with the employment plan and whether it can win the war. It is then the job of the U.S. Transportation Command (USTRANSCOM) to determine how to get the CINC's required forces, equipment, and support units into theater by the time they are needed. The war-fighting CINC requires feedback if it is determined that his employment plan is transportationally infeasible. If not enough transportation resources are available to meet the deadlines, then some replanning may need to be done on the employment end, priorities may need to be shifted, or more resources may need to be negotiated from the commercial sector. In a crisis situation, this feasibility analysis and replanning cycle needs to happen quickly, on the order of hours.

The first product of the Planning Initiative, DART (Dynamic Analysis and Replanning Tool), was built for USTRANSCOM to automate and speed up the editing and evaluation of deployment plans. These "plans" are called TPFDDs, and specify all the units being moved, including major forces, supporting forces, or resupply, the mode of transport for each unit (air, land, or sea), and the partial schedule of departure dates and arrival dates. DART supports the analysis of TPFDDs by passing the information through standard simulations and receiving feedback on late arrivals of people or cargo. Planners at USTRANSCOM are then able to use DART to highlight some of the problem areas in the TPFDD so that they may intelligently revise the TPFDD before analyzing it again. Operational users have been trained on the DART system, and they are using it to help them to their day-to-day jobs.

Currently, however, the CINC's staff has no automated support for initially generating plans for force employment or deployment; the plans are still built up by hand. The truth is that there is not a single generative planning system in operational use anywhere in the military today. The technology has so far been considered too immature for real world-sized applications. One of the goals of the Planning Initiative is to show that AI technology is ready to be integrated into decision aids for the military operations planning/replanning process. The first system to step up to the challenge is SRI International's System for Operations Crisis Action Planning (SOCAP), which incorporates the SIPE-2 planning system also developed at SRI.

The next section talks about the SOCAP system, and the rest of the paper focuses on the features of the SIPE-2 planner embedded in it that contributed to its success. Attention is also given to those aspects of SIPE-2 that caused problems, or are considered areas for future work.

2 SOCAP

SOCAP was built by SRI International to demonstrate the feasibility of applying the planning technology of SIPE-2 to the generation of large-scale military operations plans. For the purposes of a large demonstration at U.S. Central Command and the Pentagon in January 1992, SOCAP was connected to DART, with a module in between that was able to take the major forces from the SOCAP-generated plan, fill in the additional support forces, sustainment, and resupply, and generate a TPFDD. After DART was used to determine the feasibility of the plan, it was shown how changes made at a high-level in the SOCAP plan (as opposed to local changes in the TPFDD) could remedy the situation.

SOCAP embodies the domain-independent planner, SIPE-2, together with the domain knowledge necessary to apply SIPE-2 to military planning, plus a user interface designed especially for military planners which makes effective use of a map display system. Figure 1 shows the SOCAP architecture, highlighting the database inputs, the user inputs, and the outputs.

On the left are the inputs to SOCAP that would come from military databases, such as a list of enemy threats and their locations, combat forces available for the operation, and geographical location information on ports, bases, etc. The user would input the goals or missions to be accomplished, and any special constraints or assumptions that he wanted to the plan to take into consideration. In the main module in the middle of the figure is the SOCAP application layer around the SIPE-2 core.

Built into the SOCAP knowledge base are the objects of the domain, the operators which describe military operations used in a plan to achieve the goal(s), the constraints that need to be maintained, and the classes of resources needed by the plan operators. The objects and their properties (for example, the combat capabilities of forces or the capacities of ships and ports) are stored in SIPE-2's sort hierarchy, while more dynamic information is represented as predicates. Because SIPE-2 is a hierarchical planner, SOCAP has operators at all levels of abstraction, and the representation of an operator is consistent whether it is a primitive action that can be taken in the world, or a subgoal that still needs to be achieved. The size of the SOCAP knowledge base is given in [2]: 200-250 classes and objects, 15-20 properties per object, around 1200 predicates, and 50-100 plan operators.

Using all this information, SOCAP will generate a plan, either automatically, or with user interaction/guidance. The user can actually decide at each level of the planning hierarchy how much of the decision making he would like to do. It was to support this capability that a new operationally oriented interface (distinct from the SIPE-2 developer's interface) was developed. The SOCAP interface provides extra capabilities for guiding the user through planning process. If desired, SOCAP can display, at each goal in the plan, the list of operators that can be used to achieve that goal. Likewise, when choosing a particular value for a variable, such as which infantry brigade to use in an operation, the user may opt to be presented with a list of brigades that satisfy the constraints on the variable and choose one himself. At the end of each plan level, SIPE-2's plan critics are called to check that the plan decisions made so far are consistent.




Figure 1

Usually, the user views the plan in the style of SIPE-2, as a network of partially-ordered plans and goals where nodes are actions and goals and arcs are ordering links. Additionally, SOCAP added the ability to view a developing plan on a map display that shows where the major forces are located on any day of the plan. This type of display is popular with military planners who think about war plans as arrows and circles on a map rather than as a graph.

When there is no more planning to be done, SOCAP outputs, in addition to the fully expanded plan network and map-based plan, a time-phased list of the major forces involved in the operation. This is what is then fleshed out to produce a TPFDD which can be shipped to the DART system for analysis.

3 SIPE-2

At the heart of the SOCAP system is SIPE-2, the artificial intelligence planner that understands how to build plans and can reason about the effects of actions taken in the world. SIPE-2 was chosen for this demonstration because it is a domain-independent planning system that has been successfully applied to several domains already, including an extended blocks world, mobile robot planning, office building construction, travel planning, and brewery production line scheduling. SOCAP added the domain-specific knowledge of the military operations planning world to the core planning engine of SIPE-2. As has been discussed briefly in the preceding section, this involves describing objects, actions with their

preconditions and effects, constraints, and resources. Also encoded in a SIPE-2 application are deductive operators that encode which changes in the world entail other changes and some domain-specific heuristics for improving the efficiency of planning.

It is significant that these things are fairly easy to represent in the SIPE-2 formalism because there is so much domain information that needs to be built in for each new scenario. The particular way in which the knowledge is built in can affect the quality of plans produced and the speed with which they are generated. For example, everything in the world could be entered into SIPE-2 as predicates; however, it is much more efficient to store the static properties of objects, such as the speed of a C-5 aircraft, in SIPE-2's sort hierarchy. This can either be viewed as an advantage or a disadvantage depending on whether the application developer is experienced with SIPE-2, or whether the user wants to add a new operator type or a new domain constraint. Some ideas have been discussed for future work to develop a user-friendly intelligent interface tool for the purpose of assisting a non-SIPE expert in extending the domain.

SIPE-2 is a hierarchical planner, which means that it has operators and goals at various levels of abstraction, and it expands all the goals at one level of a plan before trying to solve any subgoals. Planning continues until all subgoals have been achieved by actions, or until it is found that a goal cannot be satisfied. This hierarchical decomposition of a top-level goal is generally a very natural way to think about the problem-solving process. It also makes planners more efficient since there are fewer applicable operators to consider at each level, and backtracking can be reduced significantly by working out conflicts at an earlier level before all the details get filled in.

It was relatively easy to group SOCAP's plan operators into sets corresponding to the various phases or levels of command in the planning process. In the demonstration scenario, the top-level goal is to protect the territorial integrity of a third world country against its neighboring enemies. The first level of the plan is to select a mission type, such as show-of-force or deterrence. This would be the kind of decision that the JCS would make. The second level of the plan identifies the threats and their locations and sets up multiple subgoals to

counter each of those threats. With the information about the threats, the war-fighting CINC would then decide on what kinds of forces to employ and how to employ them to counter the threats. Thus, level 3 of the SOCAP plan expands each of the previous subgoals into an employment action using a particular force or unit and a deployment subgoal for the unit to arrive at the destination of the threat in time. Actually this step should be broken up into two levels, since the CINC usually specifies only what type of force is required, and the "sourcing" or selecting of specific forces is done by FORSCOM. At this stage of the planning process, it is TRANSCOM's role to flesh out the deployment subgoals of the designated forces, along with supporting forces, and sustainment and resupply, given the available transportation assets. This is reflected in the plan that has developed after level 3, as the employment actions are all filled in and cannot be broken down further, but deployments are still shown as goals.

As you can see, the employment portion of the plan does not include much detail. But SIPE-2's hierarchical nature will make it easy to add in the lower-level employment operators in the future. As a demonstration, however, SOCAP was quite successful in exploiting the hierarchical planning capability of SIPE-2 to match the military planning process.

Besides complexity and speed (the lack thereof), one of the reasons that generative planners are not currently used in the military is that operational users want to be in charge of the planning process rather than letting a computer tell them how to conduct a military campaign. And rightfully so; humans have much valuable information and insight that is too subtle, or just too massive, for a computer to handle. But there is no reason why an intelligent computer can't share the burden.

SIPE has always been capable of both automatic and interactive planning. It can be so interactive, in fact, that SIPE's role becomes one of merely guiding the user through the process of building a plan, offering suggestions, and recording decisions made. The user can vary the level of interaction as desired during the planning process. However, SIPE-2 is so flexible about the interaction, what items can be displayed and what choices can be made, that it can be too confusing for a non-SIPE expert. SOCAP's

interface is domain-specific and extracts the information from SIPE-2 that a military planner cares about. For example, SIPE-2 hides a lot of constraint information, so when choosing a value for a variable on an operator, it might not be obvious which choices satisfy the constraints imposed by that operator. The builders of SOCAP decided to extract the constraint information right away and only present to the user the consistent choices.

SIPE-2 is also a non-linear planner, so it can represent actions that are unordered with respect to each other (and possibly simultaneous). By not forcing ordering links between actions, SIPE-2 reduces the amount of backtracking caused by an action's effects violating the preconditions of a later action. Also, for some plans you want to represent possibly simultaneous actions, where they can either be executed in parallel or it doesn't matter which is executed first. (It may be interesting to note that SIPE-2 does not have the temporal reasoning capability to make this distinction.) This non-linearity of plans was an extremely important feature for military operations plans, where several actions are being executed at the same time using sharable or completely different sets of resources. In the demonstration scenario, land, sea, and air forces are able to all work on their missions simultaneously, and many deployment actions also occur on the same day, simply using different resources.

This brings up the issue of SIPE-2's special handling of resources. SIPE was the first AI planner to allow resources for an operator to be specified explicitly. SIPE-2 has been enhanced to handle reusable and consumable resources. For example, a transport-by-sea operator, instead of having a precondition of "big enough ship available during time-period1", would have "ship" listed as a reusable resource, and SIPE-2's resource contention critics would check to make sure the resource was available at that time.

SIPE-2 also has mechanisms for handling sharable resources, but in designing SOCAP, they were found to be too inflexible. For instance, as Roberto Desimone writes in [2]:

"a large military unit, such as a division, may be employed in several operations simultaneously, where each operation uses some of the division's capabilities. The

number of operations over which the division may be shared depends of the amount of resources required for each operation. Thus, the only way to reason about the shared resource is to consider the capabilities of the division as a consumable resource purely for this specific set of operations."

Even the resource reasoning capabilities that SIPE-2 does embody were not used effectively in SOCAP because of the lack of temporal reasoning capability. SIPE-2 has no concept of how long an action takes, or of time intervals between actions. Therefore, sometimes it appears that there is a resource conflict between unordered actions when, if there was only information about the times during which a resource was unavailable or about the time intervals of the actions, there might not be a problem. In SOCAP, there are dates associated with actions in the plan, but these numbers are not used at all by SIPE-2, and one may find two operations, one with a date of 24 and one with a date of 26, that are unordered with respect to each other in the plan.

One of the big wins for SOCAP is SIPE-2's powerful constraint representation language, which is richer than that of most planning systems. The ability to post constraints on the variables of plan operators as they are added to the plan can save a lot of time over the method of choosing a binding for the variable right away. The variable binding decision can sometimes be delayed until the constraints posted and propagated to it point to a single value. Also, having declarative constraints in the system aids the user interface in generating explanations for why a particular planning decision doesn't apply.

One of the difficulties with the constraint language is that it can only handle hard constraints. In most scheduling systems, it is recognized that there may not exist a schedule that meets all the constraints, and the system will allow some constraints to be relaxed. Most planning systems, however, have only dealt with hard constraints. There is a notion of priorities or preferences here that needs to be implemented; if a plan can't satisfy all the constraints, then the least important ones should be relaxed first.

SIPE-2 also supports a context mechanism whereby a user can explore multiple plan

options concurrently. This is a necessary capability for the CINC's staff, who are generating multiple COA's for a crisis so that the best one can be selected. The context mechanism builds hierarchical trees of alternative plans just as the plans themselves are hierarchical. When one path seems to lead to failure, the user can back the system up to a previous context and try another branch.

Last, but certainly not least, in the features of SIPE-2 that are valuable for crisis action planning, is the ability to support replanning during execution given some new information about the world. A plan that has been generated by SIPE-2 retains special nodes that contain the rationale for adding an action to the plan. There are also "phantom" nodes in a plan (usually not displayed) that are reminders of preconditions that we thought would not need to be explicitly achieved because some other part of the plan ensured that they would be true. The information in these special nodes is crucial for replanning tasks where something unexpected happens, say making a precondition false that was expected to be true during execution. SIPE-2's execution monitor can accept predicates about the state of the world, and it has several mechanisms for repairing a plan under various circumstances.

The demonstration of SOCAP in January 1992 focussed mostly on plan generation, but future work will put emphasis on the replanning problem. In crisis action situations the state of the world is changing rapidly, and no one is sure if a plan will actually succeed in its execution. A goal for a future demonstration will be to feed the output of a running simulation (predicates about conditions that have been made true or false) directly to the SOCAP planner so that SIPE-2 can repair the plan on the fly and still achieve the goals of the mission.

4 Conclusion

Despite the difficulties in applying SIPE-2 to the SOCAP domain, it was a very good choice because SIPE-2 has had as its main design goals user interactivity and efficiency. Historically, other generative planners have not been as strong in these two areas, a problem

which has hindered their transition from the laboratory to the operational environment.

SOCAP has successfully demonstrated that the use of state-of-the-art AI planning technology can speed up the crisis action planning process, giving commanders more time to consider a greater number of alternative COAs before selecting one and to fully analyze an operations plan before embarking upon its execution. This is in contrast to the current situation where action sometimes must be taken in the absence of a fully developed plan.

The lasting impact of the SOCAP demonstration will be to facilitate the acceptance of generative planning technology, and hopefully artificial intelligence in general, by the military and other highly operational organizations.

Acknowledgements

This work was performed by SRI International for the DARPA/Rome Laboratory Planning Initiative under contract number F30602-91-C-0039. I wish to acknowledge Marie Bienkowski and Roberto Desimone from the SOCAP team, and David Wilkins, who developed SIPE-2, for their contributions to this paper.

References

- [1] Bienkowski, Marie, "Initial Requirements Analysis for Automation in Support of Joint Military Planning", Technical Report ITAD-2062-TR-92-40, prepared for Rome Laboratory under contract number F30602-91-C-0039, March 1992.
- [2] Desimone, Roberto, "Lessons Learned in Applying SIPE-2 to the Military Operations Crisis Action Planning Domain", AAAI Spring Symposium, Stanford, CA, March 27, 1992.
- [3] Wilkins, David and Desimone, Roberto, "Applying an AI Planner to Military Operations Planning", Draft for publication in Morgan Kaufman volume on intelligent scheduling systems, July 24, 1992.